

Wavelet Compression of MODIS Satellite Images^{*}

Andrew Lenharth
Richard E. Ladner

Department of Computer Science and Engineering

Scott Hauck
Eve A. Riskin

Agnieszka Miguel

Department of Electrical Engineering

University of Washington
Seattle, WA 98195

Abstract - Moderate Resolution Imaging Spectroradiometer (MODIS) data are collected from the Terra satellite at a rate of 11 megabits/second. These data are to be transmitted to earth, stored, and analyzed by various algorithms. One solution to reducing the bandwidth and storage need for these data is to apply wavelet-based compression, which is known to provide good performance. We present a controlled study that shows the effect of data compression on cloud cover classification of MODIS data. We find that cloudcover calculations can tolerate very large compression ratios. The correlation between frequency bands can be used to compress MODIS images even further.

1 BACKGROUND

1.1 Data compression

Data compression is becoming increasingly common in modern computer and communication systems. As processing power increases faster than bandwidth, it becomes cost effective to compress the data stream to more effectively use available communication resources.

Compression removes redundancy in data. There are many standard ways to do this, including signal transforms to transform the data into a domain where redundancy is more easily discerned; building statistical models of the data; and entropy encoding the data. For a good overview of general data compression algorithms see [1].

In recent years, the wavelet transform has been used extensively to compress image and video data. Several codecs have been developed around wavelets, including EZW [11], SPIHT [4], GTW [8] and JPEG 2000 [9]. Wavelet transforms are lossless transforms that project time series data onto wavelet basis functions, producing a time-frequency signal representation. For further reading on wavelets, an excellent introduction is [2]. A very mathematical introduction can be found in [3].

The wavelet transform of an image is performed by applying two-dimensional high pass and low pass filters to

yield a quarter-sized low resolution subimage and three quarter-sized detail coefficient subimages. The low resolution image is then recursively filtered as many times as desired. The wavelet coefficients contain all the information of the original image, but much more compactly. The original image can be recovered from the transformed image using inverse filters. There are a number of effective wavelet filters that are used in practice; one of the most popular is the Daubechies biorthogonal 9/7 filter [12]. When a wavelet transform is applied to a natural image, the energy of the image tends to be concentrated in a very small portion of the wavelet coefficients. This lends itself well to many types of subsequent encoding

Essentially all the wavelet compression algorithms use bit-plane encoding. That is, each wavelet coefficient is represented by a signed b bit number. Bit-plane encoding first encodes the most significant bit of all the coefficients, then the next most significant bit of all the coefficients, and so on. The sign of a coefficient is encoded only when it is needed (in the first bit-plane where the bit of the coefficient is a 1). Lossy compression is achieved by truncating the encoded sequence at some point. Every prefix of the encoding can be decoded into a lower quality version of the original image.

The algorithm used in this work is Set Partitioning in Hierarchical Trees (SPIHT) [4]. SPIHT has excellent rate-distortion performance and is fast as a sequential program. In addition, it can be parallelized for implementation on an FPGA (Field Programmable Gate Array), which means it has great potential for use on satellites [13].

1.2. Data Sets

The data used are from the Moderate Resolution Imaging Spectroradiometer (MODIS) instrument. MODIS is designed to collect 50 wavelengths of light, but the only satellite currently in orbit collects 36 wavelengths. The resolution of the data ranges from 250 m to 1000 m pixels, depending on which wavelengths are being recorded. Each pixel is represented in 16 bits although the precision of the sensor is only 12 bits. The data sets used in these experiments come

^{*} Research supported by NASA Contract NAS5-00213.

from the MODIS airborne simulator (MAS). The MAS flights were test flights aboard NASA U2s to test the MODIS prior to launch and to provide data from which to develop analysis algorithms. The smaller data set is from flight 96-108 and is 810x716 pixels and 50 bands. This set consists mostly of flat farmlands with some lakes and a 20% cloud cover, as can be seen in figure 2. The second set is from MAS flight 97-045 track 6, and is a mixture of snow, ice, and rock. This set has about 1% cloud cover. This set is much larger at 12,490x716 pixels and 50 bands.

1.3 Analysis algorithms:

The analysis algorithms applied to the MODIS data are usually computationally very simple because the amount of data from the satellite is so great. The analysis mostly consists of taking ratios of sums and differences between bands to form indexes to indicate the level of some property. For example, there is a vegetation index measuring the amount of vegetable matter in the area represented by the pixel, a snow index indicating how much snow is there, etc. These indexes are analyzed with simple rules to classify pixels.

A cloudmask program [5] designed for the MAS data is used to analyze the data (the implementation is the University of Wisconsin cloudmask program version release-2-0 (8/1998)). The algorithm uses 12 bands of the image to identify cloud-obstructed pixels. It also attempts to determine the type of cloud and whether a pixel is in a cloud shadow. To narrow down the information, the “Unobstructed FOV Quality Flag” is the only data field that we examine. This field is two bits, and gives 4 readings: Cloud, 66% probability clear, 95% probability clear, and 99% probability clear.

1.4 Goal of Studies:

We will study the effects of wavelet compression on the analysis of the data using several tools. Our analysis will be performed using the cloudmask algorithm. A pixel will be considered to have changed classification if and only if the “Unobstructed FOV Quality Flag” field is different between the analysis of the compressed data and the analysis of the original data. This provides a simple test for error, and is main classification result of the cloudmask algorithm. To compress the data, we use the SPIHT algorithm.

2 STUDIES

2.1 Equal Rate Individual Band Compression

The first study was designed to test the effects of wavelet compression on the results of the cloud cover analysis. First, each band was extracted from the data set. Then each band was compressed using SPIHT to the same bit rate. The data set was analyzed with the cloudmask algorithm and the

results compared with the results of the original analysis. All changes in the “Unobstructed FOV Quality Flag” classification of pixels were tracked.

The results were very satisfactory. As can be seen in figure 1, for the large data set, the error is under 5% for bit rates above 0.1 bpp. Large portions of the error are changes in certainty of 1 step, such as “cloud” to “66% probability clear”, or “95% probability clear” to “99% probability clear”. Other data sets yielded similar results.




Figure 1. Incorrect Classification in Larger Data Set

We now examine the resulting cloudmasks for the smaller data set. Figure 2 shows from left to right, the original cloudmask, the cloudmask at 0.1 bpp, and the pixels where these two images disagree. In the first two images, red represents clouds, green represents 66% probably clear, black represents 95% probably clear and white represents 99% probably clear. It is clear that the errors are concentrated around the edges. The most common artifact of wavelet compression on images is edge blurring and loss of texture. It is important to note that no blocking occurred and there are no large changes in classification away from the edges.

2.2 Cloudmask Optimized Bit Assignment

The second study was designed to optimize the accuracy of the cloudmask analysis by variable allocation of bits between bands, while still maintaining a constant average bit rate. Some of the motivation for this comes from the later experiment dealing with the PSNR of different bands. Further motivation comes from the assumption that not all bands are equally important to the cloudmask algorithm. In fact, it uses only 12 of the 50 bands, but even within those, they are not equally important.

An exhaustive search of all possible bit rate assignments would require finding the minimum of a function of 12 dimensions. Since this is computationally infeasible, several solutions were considered. A classic solution to this problem is simulated annealing. Unfortunately, each test takes

considerable time (on the order of minutes) so something simpler was needed.

We chose an average bit rate of 1 bpp and a minimum bit rate of .5 bpp. Then, setting all bands but one to the minimum bit rate, we allocated all extra bits to the one selected band. By thus restricting the degrees of freedom to 1, an optimal solution could be found under this constraint. The results show the analysis only benefited when extra bits were assigned to bands 1, 31, and 45, with band 31 giving 50% less error than average.

We then repeated the process, splitting the extra bits into two equal sized chunks, and assigning those to all pairs of bands. All combinations that performed better than average then had the bit rates varied between the two relevant bands. This allowed an efficient way to find a good solution to the problem with three degrees of freedom (the two bands, and the way the bits are divided among them). We discovered that bands 31 and 45 are by far the most important bands for the calculation of the cloudmask. The optimal assignment of bits between these bands is very close to even assignment, and gave an error of less than 1%. In contrast, equal assignment of bit rates to all bands gave an error of about 2.3%, so uneven assignment of bit rates between bands did give significantly better results for cloudmask analysis.

Unfortunately, this only tests the sensitivity of the cloudmask algorithm on this particular data set and does not give sufficient insight on the effects of compression to the analysis of these data sets. A larger variety of analysis algorithms are required to determine an optimal bit assignment.

2.3 PSNR of Each Band

Since the cloudmask algorithm was not sufficient to gauge the sensitivity of the bands to wavelet compression, a more general approach was needed. We therefore plotted the peak signal-to-noise ratio (PSNR) of each band as we varied

the bit rate.

The PSNR graph for the smaller data set is shown in figure 3. This is very similar to the PSNR graph for the larger data set, as expected. Interestingly, there are two bands in this data set, and three in the larger data set, with very low PSNR. These are likely some of the 14 bands dropped from the MODIS instrument when it was launched. The real MODIS data sets have only 36 bands.

The PSNR values may seem high for image compression, since normally values range from 20db to 40db. The reason for this is that there was sensor noise in the two bands with low PSNR. This meant the maximum value used in the PSNR calculations was 2^{15} rather than 2^{12} . The effect is that the PSNR of most of the bands is effectively shifted upwards, since the error relative to a range of values of 2^{15} is much smaller than the error relative to the actual sensor range of 2^{12} . Also, a PSNR of 140 on figure 3 indicates the image was identical to the original.




Figure 3. PSNR of larger data set




Figure 2. Resulting Cloudmasks: original, .1 bpp, pixels that differ

The PSNR graphs are consistent with the cloudmask graphs. The PSNR steadily decreases as the bit rate decreases, until about 0.2 bpp. At this point the PSNR rapidly decreases. Interestingly, the majority of bands fall within a 20 db range for a given bit rate, which is a significant amount of difference. This suggests that a good algorithm may be to compress each band to a bit rate that corresponds to a fixed PSNR, rather than compressing to a fixed bit rate.

2.4 Actual Prediction Ability of Simple Predictor

Another way to improve compression performance is to take advantage of the correlation between spectral bands by using predictive coding. If an optimal band ordering, based on minimizing the compressed size of the resultant image, could be determined, compression results would be further improved [10]. However, this is another NP-complete problem, so we instead examine how well each band can predict another, without finding the optimal ordering.

Using bzip2 we determine the compressed size of files containing the difference between bands. Bzip2 is a standard Unix compression tool utilizing the Burrow-Wheeler transform [6, 7], move-to-front, run-length, and Huffman encoding. Figure 4 shows the sizes of the band differences. The diagonal, which is near zero, is ignored in the figure to make the rest of the differences more visible.




Figure 4. Compressed size in bytes of band differences in small data set. Lighter indicates smaller compressed differences.

Given these sizes, an ordering of the input bands needs to be determined to minimize the total compressed size of all bands. This is the same minimization problem as before, and in

the general case, too computationally intensive to easily solve. We calculated some partial orderings, but not a full ordering of the bands.

3 CONCLUSION

SPIHT compression is shown to be suitable for MODIS data in that it does not introduce significant error in cloudmask analysis of these data. At high bit rates such as 6 bpp, very little error is introduced into the analysis algorithms, and still a substantial saving is achieved from the raw 16 bpp. At lower bit rates such as 1 bpp, the cloudcover classification accuracy is still very high with less than 5% error introduced as a result of the compression.

There is a high amount of redundancy between bands that can be exploited to improve compression ratios further, or allow for efficient lossless encoding. Some reordering of the bands is necessary to take full advantage of this redundancy.

Several areas remain to be studied. Most prominent is compressing the bands to an equal PSNR rather than bit rate. Also, much of the analysis needs to be replicated using algorithms besides the cloudmask to ensure the soundness of the results.

REFERENCES

- [1] The Data Compression Book, Second Edition, Mark Nelson and Jean-Loup Gailly, M&T Books, 1996
- [2] Introduction to Wavelets and Wavelet Transforms, A Primer, Burrus, Gopinath, Guo, Prentice Hall, 1998
- [3] Wavelets Made Easy, Yves Nievergelt, Birkhauser, 1999
- [4] A New Fast and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees, Amir Said and William A. Pearlman, *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, number 3, pp. 243-250, June 1996
- [5] Discriminating Clear-Sky From Clouds with MODIS, Ackerman, et al., *Algorithm Theoretical Basis Document (MOD35)*
- [6] A Block-sorting Lossless Data Compression Algorithm, M. Burrows and D.J. Wheeler, May 10, 1994
- [7] Data Compression with the Burrows-Wheeler Transform, Nelson, *Dr. Dobbs's Journal*, September 1996
- [8] Group Testing for Image Compression, E. S. Hong and R. E. Ladner, *Proceedings Data Compression Conference*, pp. 3-12, March 2000.
- [9] High Performance Image Compression with EBCOT, D. Taubman, *IEEE Transactions on Image Processing*, vol. 9, num 7, pp. 1158-1170, July 2000.

[10] Band Ordering in Lossless Compression of Multispectral Images, S. R. Tate, Proceedings Data Compression Conference, pp. 311-320, March 1994.

[11] Embedded Image Coding Using Zerotrees of Wavelet Coefficients, J. M. Shapiro, IEEE Transactions on Signal Processing, vol 41, number 12, pp. 3445-3462, December 1993.

[12] Image Coding Using Wavelet Transform, M. Antonini, M. Barlaud P. Mathieu, and I. Daubechies, IEEE Transactions on Image Processing, vol. 1, number 2, pp. 205-220, April 1992.

[13] Hyperspectral Image Compression on Reconfigurable Platforms, T. Fry, M.S. Thesis, University of Washington, Dept. of Electrical Engineering, 2001.